

Iterative non-local shrinkage algorithm for undersampled MR image reconstruction

Yasir Q. Mohsin, Gregory Ongie, and Mathews Jacob, *Senior Member, IEEE*

Abstract

We introduce a fast iterative non-local shrinkage algorithm to recover MRI data from undersampled Fourier measurements. This approach is enabled by the reformulation of current non-local schemes as an alternating algorithm to minimize a global criterion. The proposed algorithm alternates between a non-local shrinkage step and a quadratic subproblem. We derive analytical shrinkage rules for several penalties that are relevant in non-local regularization. The redundancy in the searches used to evaluate the shrinkage steps are exploited using filtering operations. The resulting algorithm is observed to be considerably faster than current alternating nonlocal algorithms. The comparisons of the proposed scheme with state-of-the-art regularization schemes show a considerable reduction in alias artifacts and preservation of edges.

Index Terms

MRI, non-local means, shrinkage, compressed sensing, denoising.

I. INTRODUCTION

Non-local means (NLM) denoising algorithms were originally introduced to exploit the similarity between patches in an image to suppress noise [1], [2], [3]. These methods recover each pixel in the denoised image as a weighted linear combination of all the pixels in the noisy image; the weights between any two pixels were estimated from the noisy image as the measure of similarity between their patch neighborhoods. This algorithm has been extended to deblurring problems by reformulating it as a regularized reconstruction scheme, where the regularization penalty is the weighted sum of square differences between all the pixel pairs in the image [4], [5], [6]. The weights are estimated from the noisy or blurred images itself, similar to classical NLM schemes. One of the difficulties in applying this scheme to challenging inverse problems (e.g. MRI recovery from under sampled data) is the dependence of the criterion on pre-specified weights; the use of the weights estimated from aliased images often preserve the alias patterns rather than suppressing them. Some authors have shown that iterating between the denoising and weight estimation step improves the quality of the images in deblurring applications [7], but often had limited success in heavily undersampled Fourier inversion problems.

The alternating NLM scheme has been recently shown to be a majorize-minimize algorithm to solve for a penalized optimization problem; the penalty term is the sum of unweighted robust distances between image patches [8], [9], [10]. The above reinterpretation was motivated by half quadratic regularization methods used in the context of pixel-based smoothness regularization [11], [12], [13], [14]. The quality of the images recovered using the resulting NLM methods are heavily dependent on the specific distance metric used for inter-patch comparisons. While convex metrics such as ℓ_1 distances may be used, nonconvex metrics that correspond to the classical NLM choices are seen to provide significantly improved results [8]. Since the direct alternation between weight estimation and optimization are not guaranteed to converge to the global minimum when nonconvex metrics are used, continuation strategies are utilized to minimize the convergence of the algorithm to local minima [8]. The main challenge associated with the implementation in [8] is the high computational complexity of the alternating minimization algorithm.

In this paper, we introduce a novel iterative algorithm to directly minimize the robust non-local criterion. This approach is based on a quadratic majorization of the patch based penalty term. Unlike the majorization used in our previous work, the weights of the quadratic terms are identical for all patch pairs, but now involves a new auxiliary variable. Similar half-quadratic strategies are widely used in the context of sparse optimization [11], [12], [13], [14]. The proposed algorithm alternates between two main steps (a) non-local shrinkage to determine the auxiliary variable, and (b) a quadratic optimization problem. We re-express the quadratic penalty involving the sum of patch differences as one involving sum of pixel differences, which enables us to solve for the quadratic sub-problem analytically. We derive analytical shrinkage expressions for a range of distance functions that are relevant for non-local regularization; this generalizes the shrinkage formulae derived by Chartand in the context of ℓ_p penalties [15]. Note that each step of the iterative shrinkage algorithm is fundamentally different from classical non-local schemes that solve an weighted quadratic optimization at each step [1], [2], [3], [4], [5], [6]. The direct evaluation of the shrinkages of the patches is computationally expensive. We propose to exploit the redundancies in the shrinkages at adjacent pixels using separable filtering operations, thereby considerably reducing the computational complexity.

Yasir Q. Mohsin and Mathews Jacob are with the Department of Electrical and Computer Engineering, Univ. Iowa, IA, USA. Gregory Ongie is with the Department of Mathematics, Univ. Iowa, IA, USA (emails: {yasir-mohsin,gregory-ongie,mathews-jacob}@uiowa.edu). This work is supported by grants NSF CCF-0844812, NSF CCF-1116067, NIH 1R21HL109710-01A1, ACS RSG-11-267-01-CCE, and ONR grant N00014-13-1-0202.

We compare the convergence of the proposed scheme with the iterative reweighted algorithm in our previous implementation [8]. We observe that the proposed scheme is approximately seven times faster than our previous iterative reweighted formulation. We also compared several distance functions in the context of iterative non-local shrinkage algorithm. Our comparisons show that saturation of the distance function is key to good performance in non-local algorithms since each patch is compared with several other patches. The saturation is needed in non-local schemes unlike local TV methods, where a specified pixel is only compared with its neighbors. Our comparisons show that the truncated $\ell_p; p = 0.5$ penalty provides the best results. We perform extensive comparisons of the scheme against local total variation (TV) regularization and a recent dictionary learning algorithm, which also exploits the similarity between image patches. The experiments demonstrates the considerable benefits in using non-local regularization.

The rest of this paper is organized as follows. We briefly describe the background in Section II. The proposed iterative non-local shrinkage algorithm is detailed in Section III, while the details of the implementation is outlined in Section IV. Section V demonstrates the performance of our method on numerous examples using CS and denoising techniques.

II. BACKGROUND

A. Unified Non-Local Formulation

The iterative algorithm that alternates between classical non-local image recovery [16] and the re-estimation of weights was shown [8] to be a majorize-minimize (MM) algorithm to solve for

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f}} \underbrace{\|\mathbf{A}\mathbf{f} - \mathbf{b}\|^2}_{\mathcal{C}(\mathbf{f})} + \lambda \mathcal{G}(\mathbf{f}), \quad (1)$$

where $\mathbf{f} \in \mathbb{C}^N$ is a vector obtained by the concatenating the rows in a 2-D image $f(\mathbf{x}), \mathbf{x} \in \mathbb{Z}^2$; $\mathbf{A} \in \mathbb{C}^{M \times N}$ is a matrix that models the measurement process; and $\mathbf{b} \in \mathbb{C}^M$ is the vector of measurements. While the first term in the cost function enforces data fidelity in k-space, the second term enforces sparsity. The regularization functional $\mathcal{G}(\mathbf{f})$ is specified by:

$$\mathcal{G}(\mathbf{f}) = \sum_{\mathbf{x}} \sum_{\mathbf{y} \in \mathbf{x} + \mathcal{N}} \varphi(P_{\mathbf{x}}(\mathbf{f}) - P_{\mathbf{y}}(\mathbf{f})). \quad (2)$$

Here, φ is an appropriately chosen potential function and $P_{\mathbf{x}}$ is a patch extraction operator which extracts an image patch centered at the spatial location \mathbf{x} from the image \mathbf{f} :

$$P_{\mathbf{x}}(\mathbf{f}) = f(\mathbf{x} + \mathbf{p}), \quad \mathbf{p} \in \mathcal{B}, \quad (3)$$

where \mathcal{B} denotes the indices in the patch. For example, if we choose a square patch of size $(2N+1)$, the set $\mathcal{B} = [-N, \dots, N] \times [-N, \dots, N]$. Similarly, \mathcal{N} are the indices of the search neighborhood; the patch $P_{\mathbf{x}}(\mathbf{f})$ is compared to all the patches whose centers are specified by $\mathbf{x} + \mathcal{N}$. For example, if we choose a square shaped neighborhood of size $2M+1$, the set $\mathcal{N} = [-M, \dots, M] \times [-M, \dots, M]$. The shape of the patches and the search neighborhood may be easily changed by re-defining the sets \mathcal{N} and \mathcal{B} .

In this paper, we focus on potential functions of the form

$$\varphi(\mathbf{g}) = \phi(\|\mathbf{g}\|), \quad (4)$$

where $\phi: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is an appropriately chosen distance metric and $\|\mathbf{g}\|^2 = \sum_{\mathbf{p} \in \mathcal{B}} |g(\mathbf{p})|^2$. However, the theory presented in this paper is general enough to work for other potential functions.

B. Solution Using Iterative Reweighted Algorithm

We showed in [8] that (1) can be solved using MM scheme, where the regularization term is majorized as the weighted sum of patch differences:

$$\mathcal{G}(\mathbf{f}) \leq \underbrace{\sum_{\mathbf{x}} \sum_{\mathbf{p} \in \mathcal{N}} w_n(\mathbf{x}, \mathbf{x} + \mathbf{p}) \|P_{\mathbf{x}}(\mathbf{f}) - P_{\mathbf{x}+\mathbf{p}}(\mathbf{f})\|^2}_{\mathcal{G}_n(\mathbf{f})}. \quad (5)$$

The weights are specified by:

$$w_n(\mathbf{x}, \mathbf{y}) = \frac{\phi'(\|P_{\mathbf{x}}(\mathbf{f}_n) - P_{\mathbf{y}}(\mathbf{f}_n)\|)}{2\|P_{\mathbf{x}}(\mathbf{f}_n) - P_{\mathbf{y}}(\mathbf{f}_n)\|}. \quad (6)$$

Here, \mathbf{f}_n is the function at the n^{th} iteration. Each iteration of the MM algorithm involves the minimization of the criterion

$$\mathbf{f}_{n+1} = \arg \min_{\mathbf{f}} \|\mathbf{A}\mathbf{f} - \mathbf{b}\|^2 + \lambda \mathcal{G}_n(\mathbf{f}). \quad (7)$$

Note that this optimization problem is essentially the classical non-local H_1 regularization scheme [16]. The alternation between (5) and the re-computation of the weights (6) will converge to the local minimum of the criterion (1). Continuation strategies

were used in [8] to improve the convergence of the algorithm to the global minimum of (1). As discussed above, one of the main challenges of the algorithm is its high computational complexity. Specifically, the conjugate gradients algorithm to solve the quadratic sub-problem converges slowly as the value of the weights increase.

III. PROPOSED ALGORITHM

A. Majorization of the Penalty Term

In this work, we will consider an alternate majorization of the potential function φ specified by

$$\varphi(\mathbf{t}) = \min_{\mathbf{s}} \left\{ \psi(\mathbf{s}) + \frac{\beta}{2} \|\mathbf{s} - \mathbf{t}\|^2 \right\}. \quad (8)$$

Here, \mathbf{s} is an auxiliary variable, $\beta > 0$ is an arbitrarily chosen scalar parameter, and ψ is a function that is dependent on φ and β . Using (8), we majorize the cost function \mathcal{C} in (1):

$$\mathcal{C}(\mathbf{f}) = \min_{\{\mathbf{s}_{\mathbf{x},\mathbf{q}}\}} \|\mathbf{A}\mathbf{f} - \mathbf{b}\|^2 + \lambda \frac{\beta}{2} \sum_{\mathbf{x},\mathbf{q}} \|P_{\mathbf{x}}(\mathbf{f}) - P_{\mathbf{x}+\mathbf{q}}(\mathbf{f}) - \mathbf{s}_{\mathbf{x},\mathbf{q}}\|^2 + \lambda \sum_{\mathbf{x}} \sum_{\mathbf{q} \in \mathcal{N}} \psi(\mathbf{s}_{\mathbf{x},\mathbf{q}}) \quad (9)$$

We use an alternating minimization algorithm to optimize (9). Specifically, we alternate between the determination of the optimal variables $\{\mathbf{s}_{\mathbf{x},\mathbf{q}}\}$, assuming \mathbf{f} to be fixed and the determination of the optimal \mathbf{f} , assuming $\{\mathbf{s}_{\mathbf{x},\mathbf{q}}\}$ to be fixed. The convergence of the above scheme could potentially be improved by including augmented Lagrangian terms [17] or Bregman iterations [18]. However, the use of these methods often interfered with the continuation strategies used to improve the convergence of the algorithm to the global minima; these schemes were designed for convex cost functions, where local minima issues do not exist.

B. The \mathbf{s} Sub-Problem: solve for $\mathbf{s}_{\mathbf{x},\mathbf{q}}$, assuming \mathbf{f} fixed

If the variable \mathbf{f} is assumed to be a constant, the determination of each of the auxiliary variables $\mathbf{s}_{\mathbf{x},\mathbf{q}}$ corresponding to different values of \mathbf{x} and \mathbf{y} can be treated independently:

$$\bar{\mathbf{s}}_{\mathbf{x},\mathbf{q}} = \arg \min_{\mathbf{s}_{\mathbf{x},\mathbf{q}}} \frac{\beta}{2} \|P_{\mathbf{x}}(\mathbf{f}) - P_{\mathbf{x}+\mathbf{q}}(\mathbf{f}) - \mathbf{s}_{\mathbf{x},\mathbf{q}}\|^2 + \psi(\mathbf{s}_{\mathbf{x},\mathbf{q}}). \quad (10)$$

We will show in the subsection III-C (see (17)) that $\mathbf{s}_{\mathbf{x},\mathbf{q}}$ can be determined analytically as a shrinkage for all penalties of interest

$$\bar{\mathbf{s}}_{\mathbf{x},\mathbf{q}} = [P_{\mathbf{x}}(\mathbf{f}) - P_{\mathbf{x}+\mathbf{q}}(\mathbf{f})] \nu(\|P_{\mathbf{x}}(\mathbf{f}) - P_{\mathbf{x}+\mathbf{q}}(\mathbf{f})\|), \quad (11)$$

where $\nu: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a function that is dependent on the distance metric ϕ .

Note that the structure of the algorithm is exactly the same for different choices of distance function; the analytical expressions for the shrinkage steps will change depending on the specific choice. We will determine the shrinkage rules corresponding to the useful penalties in the next section.

C. Determination of Shrinkage Rules

Several potential functions are currently available [8], resulting in different flavors of robust non-local regularization. The shrinkage rules for the cases for ℓ_p penalties with $p \leq 1$ are available in [15]. We now determine the corresponding shrinkage rules for a larger class of non-local penalties.

The majorization rule in (8) can be rewritten as:

$$\underbrace{\frac{\|\mathbf{t}\|^2}{2} - \frac{1}{\beta} \varphi(\mathbf{t})}_{r(\mathbf{t})} = \max_{\mathbf{s}} \left\{ \langle \mathbf{s}, \mathbf{t} \rangle - \underbrace{\left(\frac{1}{\beta} \psi(\mathbf{s}) + \frac{\|\mathbf{s}\|^2}{2} \right)}_{g(\mathbf{s})} \right\} \quad (12)$$

From the theory in [19], the above majorization relation is satisfied when r is a convex function, in which case $g = r^*$, the Legendre-Fenchel dual (or convex conjugate) of r :

$$r^*(\mathbf{s}) = \max_{\mathbf{t}} \{ \langle \mathbf{s}, \mathbf{t} \rangle - r(\mathbf{t}) \}. \quad (13)$$

However, the function r is not convex for most penalties φ that we are interested in, especially for small values of β . When r is not convex, we propose to approximate r by a convex function \hat{r} so that relation (12) is satisfied. We choose \hat{r} such that the epigraph of \hat{r} is the convex hull of the epigraph of r ; \hat{r} is thus the closest convex function to r ; see Fig (1 b). For φ functions

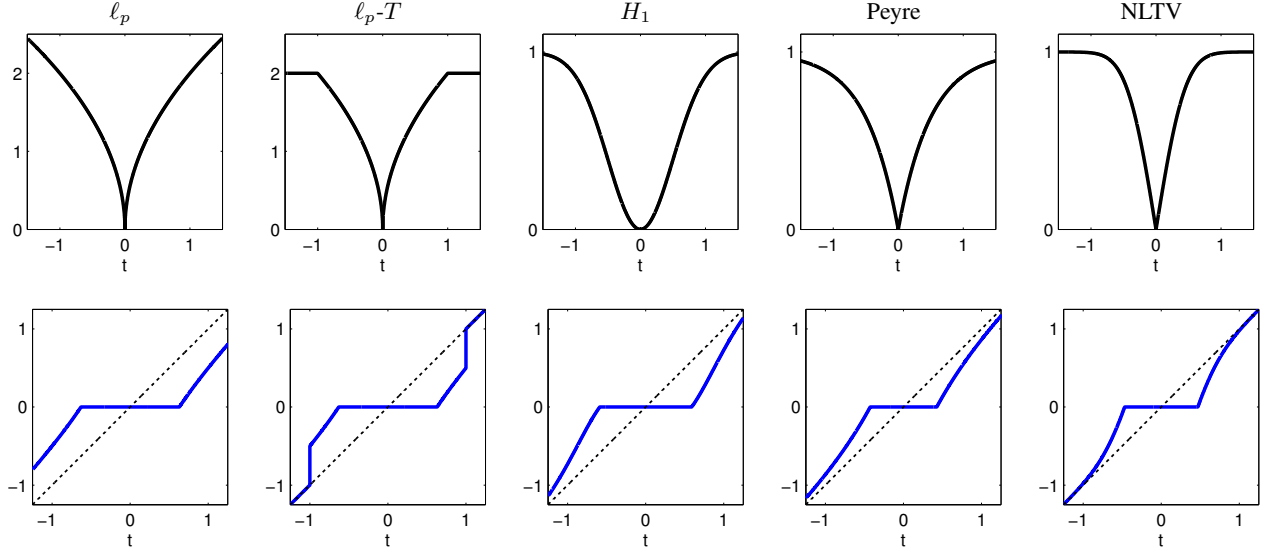


TABLE I: Distances functions $\phi(t)$ that are relevant in non-local regularization (first row) and the associated shrinkage rules $t \cdot \nu(|t|)$ (second row); see Appendix B for the corresponding formulas. Here we illustrate the shrinkage rules in 1-D for the parameter choices $\beta = 2$, $p = 0.5$, $T = 1$, and $\sigma = 0.5$. The approach introduced in the paper enables the evaluation of shrinkage rules for a much larger class of penalties, generalizing the results in [15] for ℓ_p penalties shown in the first column.

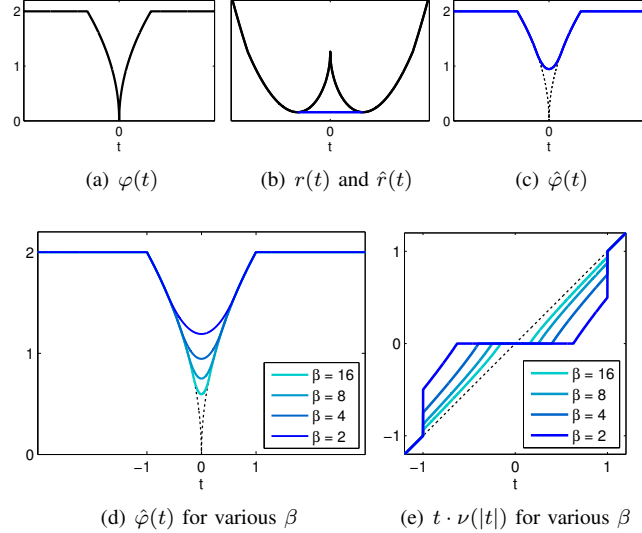


Fig. 1: Approximation of the potential function: (a) shows the original potential function $\varphi(t)$ in 1-D, which is the truncated ℓ_p ; $p = 0.5$ penalty, $T = 1$. (b) indicates the corresponding $r(t) = t^2/2 - \frac{1}{\beta}\varphi(t)$ function with $\beta = 2$, shown in black. Note that this function is non-convex. Hence, we approximate this function by $\hat{r}(t)$ shown in blue, which is the best convex approximation of $r(t)$. The corresponding modified potential function is shown in blue in (c). (d) indicates the approximations for different values of β . Note that the approximations converge uniformly to φ . (e) shows the corresponding shrinkage rules. The potential functions and shrinkage rules for different penalties are shown in Table I.

of the form (4), we have $r(\mathbf{t}) = q(\|\mathbf{t}\|)$, where the function $q : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is specified by $q(t) = t^2/2 - \phi(t)/\beta$. In all the cases we consider in this paper (see Appendix B), we can obtain the convex hull approximation of r as

$$\hat{r}(\mathbf{t}) = \begin{cases} q(\|\mathbf{t}\|) & \text{if } q'(\|\mathbf{t}\|) > 0 \\ c & \text{else} \end{cases}, \quad (14)$$

where c is an appropriately chosen constant to ensure continuity of \hat{r} ; check Fig. (1 b).

The above convex hull approximation \hat{r} of r yields a cooresponding approximation $\hat{\varphi}$ of the potential function φ given as

$$\hat{\varphi}(\mathbf{t}) := \beta \left(\frac{\|\mathbf{t}\|^2}{2} - \hat{r}(\mathbf{t}) \right);$$

see Fig. (1 c). For the potentials considered in this paper this results in the ‘‘Huber-like’’ approximation:

$$\hat{\varphi}(\mathbf{t}) = \begin{cases} \frac{\beta \|\mathbf{t}\|^2}{2} - c & \text{if } \|\mathbf{t}\| < L, \\ \phi(\|\mathbf{t}\|) & \text{else} \end{cases},$$

where $L = L(\beta) \rightarrow 0$ as $\beta \rightarrow \infty$. In particular we have $\hat{\varphi} \rightarrow \varphi$ uniformly as $\beta \rightarrow \infty$; see Fig. (1 d). Therefore, the following shrinkage rules can be interpreted as those corresponding to a Huber approximation of the potential φ , where the approximation improves with increasing β .

The shrinkage rule in (10) involves the computation of $\bar{\mathbf{s}}$ specified by:

$$\begin{aligned} \bar{\mathbf{s}}(\mathbf{t}) &= \arg \min_{\mathbf{s}} \left\{ \psi(\mathbf{s}) + \frac{\beta}{2} \|\mathbf{s} - \mathbf{t}\|^2 \right\} \\ &= \arg \max_{\mathbf{s}} \{ \langle \mathbf{s}, \mathbf{t} \rangle - \hat{r}^*(\mathbf{s}) \}, \end{aligned} \quad (15)$$

which is often called the proximal mapping of ψ . However, by exploiting duality we do not explicitly access to ψ or \hat{r}^* to compute $\bar{\mathbf{s}}$. Differentiating the right hand side of (15) with respect to \mathbf{s} and setting it to zero, we obtain $\mathbf{t} - \partial \hat{r}^*(\bar{\mathbf{s}}) \ni 0$, or equivalently, $\bar{\mathbf{s}}(\mathbf{t}) \in (\partial \hat{r}^*)^{-1}(\mathbf{t})$. Since the subgradients of Legendre-Fenchel duals satisfy $(\partial \hat{r}^*)^{-1}(\mathbf{t}) = (\partial \hat{r})(\mathbf{t})$, we have

$$\bar{\mathbf{s}}(\mathbf{t}) \in \partial \hat{r}(\mathbf{t}). \quad (16)$$

Considering the expression for the convex hull approximation of \hat{r} in (14), we have:

$$\bar{\mathbf{s}}(\mathbf{t}) = \begin{cases} q'(\|\mathbf{t}\|) \frac{\mathbf{t}}{\|\mathbf{t}\|} & \text{if } q'(\|\mathbf{t}\|) > 0 \\ 0 & \text{else.} \end{cases}$$

Setting $q(t) = t^2/2 - \phi(t)/\beta$ in the above equation, we obtain the shrinkage rule as:

$$\begin{aligned} \bar{\mathbf{s}}(\mathbf{t}) &= \frac{\mathbf{t}}{\|\mathbf{t}\|} \left(\|\mathbf{t}\| - \frac{1}{\beta} \phi'(\|\mathbf{t}\|) \right)_+ \\ &= \underbrace{\mathbf{t} \left(1 - \frac{\phi'(\|\mathbf{t}\|)}{\beta \|\mathbf{t}\|} \right)_+}_{\nu(\|\mathbf{t}\|)}, \end{aligned} \quad (17)$$

where $(\cdot)_+ := \max\{\cdot, 0\}$. Here, $\nu(\|\mathbf{t}\|)$ is a scalar between 0 and 1, which when multiplied by \mathbf{t} will yield the shrinkage of \mathbf{t} . Setting $\mathbf{t} = P_{\mathbf{x}}(\mathbf{f}) - P_{\mathbf{x}+\mathbf{q}}(\mathbf{f})$ in the above equation, we obtain the shrinkage rules to be used in (11). Note that the above approach can be adapted to most penalties. We determine the shrinkage rules and the associated ν functions for common penalty functions φ in non-local regularization in Appendix B. Table I shows the penalty functions for different metrics and the corresponding shrinkage weights.

D. The f Sub-Problem: solve for f , assuming $\mathbf{s}_{\mathbf{x},\mathbf{q}}$ fixed

In this step, we assume the auxiliary variables $\mathbf{s}_{\mathbf{x},\mathbf{q}}$ to be fixed. Hence, the minimization of (1) simplifies to:

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f}} \|\mathbf{A}\mathbf{f} - \mathbf{b}\|^2 + \frac{\lambda\beta}{2} \underbrace{\sum_{\mathbf{x},\mathbf{q} \in \mathcal{N}} \|P_{\mathbf{x}}(\mathbf{f}) - P_{\mathbf{x}+\mathbf{q}}(\mathbf{f}) - \mathbf{s}_{\mathbf{x},\mathbf{q}}\|^2}_{\mathcal{R}(\mathbf{f})} \quad (18)$$

The above quadratic penalty may be solved using the conjugate gradients algorithm. However, we will now simplify it to an expression that can be solved analytically, which is considerably more efficient.

The quadratic penalty term involves differences between multiple patches in the image, each of which is a linear combination of quadratic differences between image pixels. The differences between two specific pixels are thus involved in different patch differences. We show in Appendix A that the pixel differences from several patches can be combined to obtain the following pixel-based penalty:

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f}} \|\mathbf{A}\mathbf{f} - \mathbf{b}\|^2 + \lambda \frac{\beta}{2} \sum_{\mathbf{q} \in \mathcal{N}} \|\mathbf{D}_{\mathbf{q}}\mathbf{f} - \mathbf{h}_{\mathbf{q}}\|^2. \quad (19)$$

Here, $\mathbf{D}_{\mathbf{q}}$ is the finite difference operator

$$(\mathbf{D}_{\mathbf{q}}\mathbf{f})(\mathbf{x}) = \mathbf{f}(\mathbf{x}) - \mathbf{f}(\mathbf{x} + \mathbf{q}). \quad (20)$$

The images $\mathbf{h}_{\mathbf{q}}(\mathbf{x})$, $\mathbf{q} \in \mathcal{N}$, are obtained by shrinking the finite difference terms $\mathbf{D}_{\mathbf{q}}\mathbf{f}$:

$$\mathbf{h}_{\mathbf{q}} = (\mathbf{D}_{\mathbf{q}}\mathbf{f}) \bullet \mathbf{v}_{\mathbf{q}}, \quad (21)$$

where \bullet denotes the entrywise multiplication of the vectors, and the pixel shrinkage weights $\mathbf{v}_{\mathbf{q}}$ for a specified spatial location \mathbf{x} are obtained by the sum of the shrinkage weights for the nearby patch pairs

$$\mathbf{v}_{\mathbf{q}}(\mathbf{x}) = \sum_{\mathbf{p} \in \mathcal{B}} \underbrace{\nu(\|P_{\mathbf{x}+\mathbf{p}}(\mathbf{f}) - P_{\mathbf{x}+\mathbf{p}+\mathbf{q}}(\mathbf{f})\|)}_{\mathbf{u}_{\mathbf{q}}(\mathbf{x})}. \quad (22)$$

Here, ν is specified by (17). We solve (19) in the Fourier domain for measurement operators \mathbf{A} that are diagonalizable in the Fourier domain, as shown in the next section.

IV. IMPLEMENTATION

We now focus on the implementation of the sub-problems. Specifically, we show that all of the above steps can be solved analytically for most penalties and measurement operators of practical interest. This enables us to realize a computationally efficient algorithm. We also introduce a continuation scheme to improve the convergence of the algorithm.

A. Analytical Solution of (18) in the Fourier Domain

The Euler-Lagrange equation for (19) is given by:

$$\left(2\mathbf{A}^H \mathbf{A} + \lambda\beta \sum_{\mathbf{q} \in \mathcal{N}} \mathbf{D}_{\mathbf{q}}^H \mathbf{D}_{\mathbf{q}} \right) \mathbf{f} = 2\mathbf{A}^H \mathbf{b} + \lambda\beta \sum_{\mathbf{q} \in \mathcal{N}} \mathbf{D}_{\mathbf{q}}^H \mathbf{h}_{\mathbf{q}} \quad (23)$$

Here \mathbf{B}^H denotes the Hermitian transpose of matrix \mathbf{B} . Note that the variables in the left hand side of (23) are fixed. Thus, this step involves the solution to a linear system of equations. In many inverse problems of interest (e.g. Fourier sampling, deblurring), the measurement operator \mathbf{A} is diagonalizable in the Fourier domain, in which case we may write $\mathbf{A}^H \mathbf{A}$ as a pointwise multiplication in the Fourier domain. For instance, in the particular case when \mathbf{A} is a Cartesian Fourier undersampling operator, we may write

$$\mathbf{A}^H \mathbf{A} \mathbf{f} = \mathcal{F}^{-1}(\mathbf{a} \bullet \mathcal{F}(\mathbf{f})) \quad (24)$$

where \mathcal{F} discrete Fourier transform and \mathbf{a} is a vector of ones and zeros corresponding to the Fourier sample locations. Likewise, assuming circular boundary conditions for the finite difference operator $\mathbf{D}_{\mathbf{q}}$, the operator $\mathbf{D}_{\mathbf{q}}^H \mathbf{D}_{\mathbf{q}}$ is diagonalizable in the Fourier domain as

$$\mathbf{D}_{\mathbf{q}}^H \mathbf{D}_{\mathbf{q}} \mathbf{f} = \mathcal{F}^{-1}(|\mathbf{d}_{\mathbf{q}}|^2 \bullet \mathcal{F}(\mathbf{f})), \quad (25)$$

where $|\mathbf{d}_{\mathbf{q}}|^2$ is the pointwise modulus squared of the Fourier multiplier $\mathbf{d}_{\mathbf{q}}$ corresponding to $\mathbf{D}_{\mathbf{q}}$. Hence, taking the DFT of both sides of (23) we have

$$\left(2\mathbf{a} + \lambda\beta \sum_{\mathbf{q} \in \mathcal{N}} |\mathbf{d}_{\mathbf{q}}|^2 \right) \bullet \mathcal{F}(\mathbf{f}) = 2\mathbf{b}_0 + \lambda\beta \mathcal{F} \left(\sum_{\mathbf{q} \in \mathcal{N}} \mathbf{D}_{\mathbf{q}}^H \mathbf{h}_{\mathbf{q}} \right),$$

where $\mathbf{b}_0 = \mathcal{F}(\mathbf{A}^H \mathbf{b}) \in \mathbb{C}^M$ is a zero-padded version of the Fourier samples $\mathbf{b} \in \mathbb{C}^N$. Solving for \mathbf{f} gives

$$\mathbf{f} = \mathcal{F}^{-1} \left[\frac{2\mathbf{b}_0 + \lambda\beta \mathcal{F} \left(\sum_{\mathbf{q} \in \mathcal{N}} \mathbf{D}_{\mathbf{q}}^H \mathbf{h}_{\mathbf{q}} \right)}{2\mathbf{a} + \lambda\beta \sum_{\mathbf{q} \in \mathcal{N}} |\mathbf{d}_{\mathbf{q}}|^2} \right] \quad (26)$$

where the division occurs entrywise.

In inverse problems such as non-Cartesian MRI and parallel MRI, where the measurement operator \mathbf{A} is not diagonalizable in the Fourier domain, we propose to solve (23) efficiently using pre-conditioned conjugate gradient (CG) algorithm. The above simplifications of the derivative operator can be used to develop an efficient pre-conditioner in these cases. A few CG steps at each iteration are often sufficient for good convergence since the algorithm is initialized by the previous iterate.

B. Efficient Evaluation of Shrinkage Weights

We now focus on the efficient evaluation of $\mathbf{v}_{\mathbf{q}}(\mathbf{x})$; $\forall \mathbf{q} \in \mathcal{N}$ in (22). Note that $\mathbf{u}_{\mathbf{q}}(\mathbf{x})$ involves the comparison of the patches $P_{\mathbf{x}}(\mathbf{f})$ and $P_{\mathbf{x}+\mathbf{q}}(\mathbf{f})$; since these quantities have to be computed for all spatial locations \mathbf{x} and different shifts \mathbf{q} , the

direct evaluation of (22) is computationally expensive. We propose to exploit the redundancies between $\mathbf{v}_q(\mathbf{x})$ to considerably accelerate their computation. From (22), we have

$$\begin{aligned} \mathbf{u}_q(\mathbf{x}) &= \nu (\|P_{\mathbf{x}}(\mathbf{f}) - P_{\mathbf{x}+\mathbf{q}}(\mathbf{f})\|) \\ &= \nu \left(\sqrt{\underbrace{\sum_{p \in \mathcal{B}} \|\mathbf{f}(\mathbf{x} - \mathbf{p}) - \mathbf{f}(\mathbf{x} - \mathbf{p} + \mathbf{q})\|^2}_{(|\mathbf{D}_q \mathbf{f}|^2 * \eta)(\mathbf{x})}} \right). \end{aligned}$$

Here η is a moving average filter with the size of the patch. The above equation implies that the computation of $\mathbf{u}_q(\mathbf{x}); \forall \mathbf{x}$ can be obtained efficiently by simple pointwise operations and a computationally efficient filtering operation. Combining the above result with (22), we obtain

$$\mathbf{v}_q = \underbrace{\left[\nu \left(\sqrt{|\mathbf{D}_q \mathbf{f}|^2 * \eta} \right) \right]}_{\mathbf{u}_q} * \eta \quad (27)$$

We realize the convolutions $|\mathbf{D}_q \mathbf{f}|^2 * \eta$ and $\mathbf{u}_q * \eta$ using separable moving average filtering operations.

C. Continuation Strategy to Improve Convergence

The quality of the majorization in (8) depends on the parameter β . It is known that high values of β results in poor convergence. However, since we require the convex-hull approximation of the original penalty (see Section III-C) for the majorization, the solution of the proposed scheme corresponds to that of the original problem only when $\beta \rightarrow \infty$. We hence use a continuation strategy to improve the convergence rate, where β is initialized with a small value and is increased gradually to a high value. This approach is adopted by us [8], [10], as well as other authors [20], [21]. We also use continuation to truncate the metric penalties in which we start with a large threshold and gradually decrease it until it attains a small value; this means that we are not concerned with the distant patches that are high possible to be dissimilar. The saturated ℓ_p norm seems to give better results than non truncated one.

The pseudo-code of the algorithm is shown below.

Algorithm IV.1: NONLOCAL SHRINKAGE($\mathbf{A}, \mathbf{b}, \lambda$)

```

Input :  $\mathbf{b}$  = k-space measurements
 $\beta = \beta_{\text{init}}; T = T_{\text{init}};$ 
while  $i < \#$  Outer Iterations
    while  $j < \#$  Inner Iterations
        do {
            Compute  $\mathbf{v}_q; \forall \mathbf{q} \in \mathcal{N}$  using (27)
            Compute  $\mathbf{h}_q; \forall \mathbf{q} \in \mathcal{N}$  using (21)
            Update  $\mathbf{f}$  according to (26)
             $\beta \leftarrow \beta * \beta_{\text{incfactor}}$ 
             $T \leftarrow T * T_{\text{decfactor}}$ 
        }
    return ( $\mathbf{f}$ )

```

We observe from the pseudo-code that the algorithm requires two moving-average filtering operations per \mathbf{q} value to evaluate (27). For a 3×3 neighborhood, this translates to 16 moving filtering operations. In addition, the evaluation of \mathbf{f} according to (26) requires one FFT and one IFFT. We typically need 20 inner iterations and about 30–40 outer iterations for the best convergence and recovery.

The algorithm was implemented in MATLAB 2012 using Jacket [22] on a Linux workstation machine with eight cores and a NVIDIA Tesla graphical processing unit. In all the experiments, we initialize $\beta = 0.01$ and update it by a factor of 2 in each outer iteration.

V. RESULTS

We will now focus on the implementation of our scheme in the contexts of CS and denoising. Some of the MR images used in these experiments are courtesy of American Radiology Services¹.

¹ www3.americanradiology.com/pls/web1/wwimggal.vmg

A. Convergence Rate

We first compare the proposed scheme with our previous iterative reweighted non-local algorithm [8]. We consider the recovery of the 256×256 MR brain image using a five fold under sampled random sampling pattern. The regularization parameters of both algorithms were set to $\lambda = 10^{-4}$; this parameter was chosen to obtain the best possible reconstruction by comparing with the original image. The convergence plots of the algorithm as a function of the CPU time are shown in Fig (2). We observe that both the algorithms converge to almost the same result. However, the non-local shrinkage algorithm converged around ten times faster than the iterative reweighted scheme. The reconstructions demonstrates the quality improvement offered by the proposed scheme for a specified computation time. One of the reasons for the faster convergence of the proposed algorithm can be attributed to the fast inversion of the quadratic sub-problems. The condition number of the quadratic subproblem in iterative reweighting [8] grows with iterations, resulting in slow convergence of the CG algorithms that were used to solve it.




Fig. 2: Comparison of the convergence of the iterative reweighted non-local algorithm and the proposed iterative non-local shrinkage algorithm. The plots indicate the evolution of the cost function specified by (1) and the signal to noise ratio (SNR) as a function of the computational time. We observe that the proposed scheme converges around ten times faster than the iterative reweighted algorithm.

B. Impact of the Distance Metric

The proposed scheme can be adapted to most non-local distance metrics by simply changing the shrinkage rule. The shrinkage rules for different non-local penalties are shown in Table I. We compare the different metrics in the context of recovering images from 5 fold acceleration using randomly under sampled data. The SNR of the reconstructions are shown in Table II. The regularization parameters of each of the algorithms are optimized to provide the best possible results. The first column corresponds to the convex ℓ_1 differences between patches. The second and third columns correspond to alternating H_1 and NLTV penalties [8], respectively. These penalties depend on the parameter σ corresponding to the width of the Gaussian weight function. This parameter is analogous to the threshold T used in the saturating ℓ_1 and ℓ_p penalties, shown in the last two columns. All of these parameters were optimized to ensure fair comparisons. We also used continuation schemes on these parameters to improve the convergence to global minima.

All of the penalties, except the ℓ_1 distance function saturate with inter-patch distances. This explains the poor performance of the convex ℓ_1 penalty compared to the non-convex counterparts. Unlike local total variation scheme, which only compares a particular pixel with a few other pixels, several pixel comparisons are involved in non-local regularization. Saturating priors are needed to avoid the averaging of dissimilar patches, which may result in blurring. Since the saturating ℓ_p metric provides the best over all reconstructions, we use this prior for all the future comparisons in this paper.

Image	ℓ_1	H_1	NLTV	ℓ_1 -T	ℓ_p -T
Brain1	22.12	27.15	28.06	27.53	28.41
Brain2	22.48	26.12	27.15	26.35	27.93
ankle	22.23	23.43	24.56	23.81	24.67

TABLE II: Impact of the distance metric on the reconstructions. We compare the reconstructions obtained using the non-local shrinkage algorithm using ℓ_1 , H_1 , NLTV, thresholded ℓ_1 and thresholded ℓ_p ; $p = 0.5$ metrics. All the metrics, except the convex ℓ_1 scheme are saturating priors. We observe that saturation is key to good performance of non-local algorithms. Among the different metrics, the thresholded ℓ_p penalty is observed to provide the best results in all the examples.

C. Comparisons With State-of-the-Art Algorithms

We compare the proposed scheme with classical local total variation algorithm and the dictionary learning MRI (DLMRI) scheme [23]. The DLMRI method is also a patch based regularization scheme, where a dictionary is learned from the patches in the image. This scheme was reported [23] to provide considerably better reconstructions than the sparse recovery algorithm

Image	DLMRI		TV		NLS	
	SNR	PSNR	SNR	PSNR	SNR	PSNR
Brain1	20.46	29.73	22.80	32.87	28.41	39.14
Brain2	21.40	31.77	23.61	34.66	27.93	39.54
Brain3	23.24	36.96	26.70	41.20	29.13	44.15
Ankle	20.10	32.15	22.22	34.83	24.67	38.00
Spine	26.53	36.76	30.39	41.52	33.11	44.73
Willis'	21.33	33.13	21.42	33.95	23.81	36.91

TABLE III: Quantitative comparison of the proposed iterative non-local shrinkage (NLS) algorithm using the saturating ℓ_p ; $p = 0.5$ penalty with dictionary learning MRI (DLMRI) [23] and local total variation regularization schemes in the absence of noise. We considered random sampling. The SNR and PSNR metrics of the reconstructed images are shown in the table.

combining wavelet and TV regularization [24]. We relied on the MATLAB implementation of DLMRI². A key difference with the results reported in [23] is that we used the complex version of the code distributed by the authors. This was required to make the comparisons fair to TV and our non-local implementations as both of them do not use this constraint. Note that this assumption is often not satisfied in routine MRI exams.

The comparison of the methods in the context of random sampling with 20% of the samples retained in the absence of noise is shown in Fig. 4. The regularization parameters of all the algorithms have been optimized to yield the best signal to noise ratio (SNR). The SNR and the peak SNR (PSNR) that are used for the comparisons in this paper are computed as

$$\text{SNR} = 20 \log_{10} \left(\frac{\|\mathbf{\Gamma}_{\text{orig}}\|_F}{\|\mathbf{\Gamma}_{\text{rec}} - \mathbf{\Gamma}_{\text{orig}}\|_F} \right)$$

$$\text{PSNR} = 20 \log_{10} \left(\frac{MAX \cdot \sqrt{N_x \cdot N_y}}{\|\mathbf{\Gamma}_{\text{rec}} - \mathbf{\Gamma}_{\text{orig}}\|_F} \right)$$

where N_x and N_y are the image dimensions, and MAX is the maximum allowed pixel intensity.

We observe that the proposed non-local algorithm provides better preservation of edge details and minimize patchy artifacts as seen in TV reconstructions. The quantitative comparisons of different methods on more MR images in the absence of noise using 5 fold random sampling operator are reported in Table III. We observe that the NLS scheme provides a consistent 2-4 dB improvement over the other methods in most cases.

D. Performance With Noise

We study the performance of the proposed algorithms as a function of acceleration in the presence of noise in Fig. 3. We used a 512×512 MRI brain image, sampled using a random sampling operator at different acceleration factors ($R = 2.5, 4, 6, 8, 10$ and 20). The measurements were contaminated with complex white Gaussian noise of $\sigma = 10.2$. The PSNR and SNR as a function of accelerations of this experiment are plotted in Fig. 3, where we compare our method against DLMRI and TV. We observe that the proposed scheme provides a consistent improvement in the presence of noise.

The reconstructions of an ankle image from its 4 fold Cartesian undersampled Fourier data, corrupted with zero mean complex Gaussian noise with a standard deviation $\sigma = 10$, are shown in Fig. 5. This is a really challenging case since the 1-D downsampling pattern is considerably less efficient than the 2-D random pattern used in the previous experiment. We observe that the non-local algorithm provides better reconstructions than the other schemes. Specifically, the TV scheme results in patchy artifacts. The DLMRI scheme results in blurring and loss of details close to the heel. The details are relatively better preserved close to the finger since there are no structures above or below it that aliases to it. By contrast to the classical algorithms, the degradation in performance of the non-local algorithm is comparatively small. The quantitative comparisons of the algorithms on this setting using different images are shown in the top section of Table IV.

The reconstructions of a 256×256 brain image from its radial samples acquired with a 40 spoke trajectory are shown in Fig. 6. The measurements are corrupted with zero mean complex Gaussian noise of standard deviation $\sigma = 18.8$. All methods result in loss of subtle image features since the acceleration factor and the noise level are high. We observe that the NLS scheme provides better recovery than the competing methods. The quantitative results in this setting for various MR images are shown in the bottom section of Table IV. We observe that the SNR improvement offered by NLS over the other methods are not as high as in the previous cases, mainly due to the considerable noise in the data and the high acceleration.

Finally, we show the recovery of four MR images from three fold radial under sampled data that is contaminated with zero mean complex Gaussian noise of standard deviation $\sigma = 10$. These experiments show that the NLS scheme can be used to obtain good quality reconstructions at moderate acceleration factors and noise levels.

²The DLMRI code is available on the author's website www.ifp.illinois.edu/~yoram/DLMRI-Lab/DLMRI.html




Fig. 3: SNR and PSNR vs Acceleration. We used a 512×512 MRI brain image, sampled using a random sampling operator at different acceleration factors ($R=2.5, 4, 6, 8, 10$ and 20). The measurements were contaminated with complex white Gaussian noise of $\sigma = 10.2$. The SNR of the reconstructions obtained using the three algorithms are plotted. These results show that the NLS scheme is capable of providing better reconstructions at a range of accelerations. Specifically, it provides a consistent improvement of 0.5-9 dB in SNR/PSNR over the other algorithms.

Image	DLMRI		TV		NLS	
	SNR	PSNR	SNR	PSNR	SNR	PSNR
Brain1	13.55	22.82	14.81	24.67	18.29	28.45
Brain2	14.38	24.74	16.10	27.12	18.63	29.83
Brain3	13.10	26.82	16.19	30.37	19.73	33.80
Ankle	12.96	25.00	15.02	27.64	18.52	31.13
Spine	16.33	26.57	18.38	29.29	20.49	31.57
Willis'	14.56	26.35	16.08	28.53	18.14	30.45
Brain1	12.59	21.87	11.84	21.27	12.35	21.81
Brain2	17.46	27.83	17.43	28.14	18.46	29.54
Brain3	14.13	27.85	16.98	31.00	18.00	31.97
Ankle	15.80	27.85	16.17	28.57	16.81	29.57
Spine	18.54	28.77	19.86	30.50	20.30	30.82
Willis'	14.18	25.97	14.55	26.69	15.61	27.69

TABLE IV: Quantitative comparison of the algorithms in the presence of noise. The top part shows the SNR of the reconstructions obtained from 4 fold Cartesian under sampled data, contaminated by zero mean complex Gaussian noise with standard deviation $\sigma = 10$. The bottom part shows the SNR of the reconstructions from radial under sampled data with 40 spokes, contaminated by zero mean complex Gaussian noise with standard deviation $\sigma = 18.8$. The quantitative results show that the proposed iterative NLS scheme provides consistently better reconstructions for the above cases.

VI. CONCLUSION

We introduced a fast iterative non-local shrinkage algorithm to recover MR image data from under sampled Fourier measurements. This approach is enabled by the reformulation of current non-local schemes as an iterative re-weighting algorithm to minimize a global criterion [8]. The proposed algorithm alternates between a non-local shrinkage step and a quadratic subproblem, which can be solved analytically and efficiently. We derived analytical shrinkage rules for several penalties that are relevant in non-local regularization. We accelerated the non-local shrinkage step, whose direct evaluation involves expensive non-local patch comparisons, by exploiting the redundancy between the terms at adjacent pixels. The resulting algorithm is observed to be considerably faster than our previous implementation. The comparison of different penalties demonstrated the benefit in using distance functions that saturate with distant patches. The comparisons of the proposed scheme with state of the art algorithms show a considerable reduction in alias artifacts and preservation of edges.

APPENDIX A: SIMPLIFICATION OF EQ (18)

Using the formula for the shrinkage from (11), specified by $\mathbf{s}_{\mathbf{x},\mathbf{q}} = (P_{\mathbf{x}}(\mathbf{f}) - P_{\mathbf{x}+\mathbf{q}}(\mathbf{f})) \nu(\|\mathbf{P}_{\mathbf{x}}(\mathbf{f}) - \mathbf{P}_{\mathbf{x}+\mathbf{q}}(\mathbf{f})\|)$, we obtain

$$\mathcal{R}(\mathbf{f}) = \sum_{\mathbf{x}} \sum_{\mathbf{q}} \|P_{\mathbf{x}}(\mathbf{f}) - P_{\mathbf{x}+\mathbf{q}}(\mathbf{f}) - \mathbf{s}_{\mathbf{x},\mathbf{q}}\|^2 \quad (28)$$

Expanding the above expression

$$\mathcal{R}(\mathbf{f}) = \sum_{\mathbf{x}} \sum_{\mathbf{q}} \sum_{\mathbf{p} \in \mathcal{B}} |f(\mathbf{x} + \mathbf{p}) - f(\mathbf{x} + \mathbf{q} + \mathbf{p}) - \mathbf{s}_{\mathbf{x},\mathbf{q}}(\mathbf{p})|^2 \quad (29)$$




Fig. 4: Comparison of the algorithms in the absence of noise. We consider the recovery of a 256×256 MRI brain image from 20% of its Fourier samples, acquired using a random sampling pattern shown in (e) using non-local shrinkage scheme (NLS), DLMRI and local TV (TV). The reconstructions are shown in (b)-(d). The corresponding error images, scaled by a factor of 5 for better visualization, are shown in the bottom row. The reconstructions show that the NLS scheme is capable of better preserving the edges and details, resulting in less blurred reconstructions.




Fig. 5: Comparison of the algorithms in the presence of noise. We consider the recovery of a 256×256 MRI ankle image from its Cartesian Fourier sampling pattern (shown in (e)), contaminated by zero mean complex Gaussian noise with standard deviation $\sigma = 10$. The top row shows the original and reconstructed images, while the error images scale by a factor of five are shown in the bottom row. This is a challenging case due to the high 1-D undersampling factors and noise. We observe that the NLS scheme provides the best reconstructions with minimal alias artifacts.

We use a change of variables $\mathbf{x} = \mathbf{x} + \mathbf{p}$ to obtain

$$\begin{aligned}\mathcal{R}(f) &= \sum_{\mathbf{x}} \sum_{\mathbf{q}} \sum_{\mathbf{p} \in \mathcal{B}} \overbrace{|f(\mathbf{x}) - f(\mathbf{x} + \mathbf{q}) - \mathbf{s}_{\mathbf{x}-\mathbf{p},\mathbf{q}}(\mathbf{p})|^2}^{e(\mathbf{x},\mathbf{q})} \\ &= \sum_{\mathbf{x},\mathbf{y}} |e(\mathbf{x},\mathbf{q})|^2 + c\end{aligned}\quad (30)$$

$$\begin{aligned}&+ 2 \sum_{\mathbf{x},\mathbf{y}} e(\mathbf{x},\mathbf{q}) \underbrace{\sum_{\mathbf{p} \in \mathcal{B}} \mathbf{s}_{\mathbf{x}-\mathbf{p},\mathbf{q}}(\mathbf{p})}_{h_{\mathbf{q}}(\mathbf{x})} \\ &= \sum_{\mathbf{x}} \sum_{\mathbf{y} \in \mathbf{x} + \mathcal{N}} |f(\mathbf{x}) - f(\mathbf{y}) - h_{\mathbf{q}}(\mathbf{x})|^2 + c - d.\end{aligned}\quad (31)$$

In the above equations, c and d are constants specified by

$$\begin{aligned}c &= \sum_{\mathbf{x},\mathbf{y}} \sum_{\mathbf{p} \in \mathcal{B}} |\mathbf{s}_{\mathbf{x}-\mathbf{p},\mathbf{q}}(\mathbf{p})|^2 \\ d &= \sum_{\mathbf{x},\mathbf{y}} |h_{\mathbf{q}}(\mathbf{x})|^2.\end{aligned}$$

Since the solution to (18) does not depend on the constants, we ignore these terms. Thus, (29) can be rewritten using (31) as

$$\mathcal{R}(f) = \sum_{\mathbf{q} \in \mathcal{N}} \underbrace{\|f(\mathbf{x}) - f(\mathbf{x} + \mathbf{q}) - h_{\mathbf{q}}(\mathbf{x})\|^2}_{\mathcal{D}_{\mathbf{q}}f}$$

Here, $\mathcal{D}_{\mathbf{q}}f(\mathbf{x}) = f(\mathbf{x} + \mathbf{q}) - f(\mathbf{x})$ is the finite difference operator.

We observe that the expression for $h_{\mathbf{q}}(\mathbf{x})$

$$h_{\mathbf{q}}(\mathbf{x}) = \sum_{\mathbf{p} \in \mathcal{B}} \mathbf{s}_{\mathbf{x}-\mathbf{p},\mathbf{q}}(\mathbf{p}), \quad (32)$$

can be further simplified. From (11), we have the patch \mathbf{s} specified as

$$\mathbf{s}_{\mathbf{x},\mathbf{q}} = (P_{\mathbf{x}}f - P_{\mathbf{x}+\mathbf{q}}f) \underbrace{\nu(\|P_{\mathbf{x}}f - P_{\mathbf{x}+\mathbf{q}}f\|)}_{u_{\mathbf{q}}(\mathbf{x})}$$

Here, $u_{\mathbf{q}}(\mathbf{x}) = \nu(\|P_{\mathbf{x}}f - P_{\mathbf{x}+\mathbf{q}}f\|)$ is the factor between 0 and 1, which is multiplied by the patch to get the shrunk patch. Hence, $\mathbf{s}_{\mathbf{x},\mathbf{q}}(\mathbf{r}) = [f(\mathbf{x} + \mathbf{r}) - f(\mathbf{x} + \mathbf{q} + \mathbf{r})] \cdot u_{\mathbf{q}}(\mathbf{x})$; $\mathbf{r} \in \mathcal{B}$. Thus, we have

$$\mathbf{s}_{\mathbf{x}-\mathbf{p},\mathbf{q}}(\mathbf{p}) = [f(\mathbf{x}) - f(\mathbf{x} + \mathbf{q})] u_{\mathbf{q}}(\mathbf{x} - \mathbf{p}). \quad (33)$$

Substituting in (32), we get

$$h_{\mathbf{q}}(\mathbf{x}) = (f(\mathbf{x}) - f(\mathbf{x} + \mathbf{q})) \underbrace{\sum_{\mathbf{p} \in \mathcal{B}} u_{\mathbf{q}}(\mathbf{x} - \mathbf{p})}_{v_{\mathbf{q}}(\mathbf{x})}. \quad (34)$$

APPENDIX B: SHRINKAGE RULES FOR USEFUL NON-LOCAL DISTANCE FUNCTIONS

1) *Thresholded ℓ_p ; $p \leq 1$ metric:* We now consider the saturating ℓ_p metric, specified by

$$\phi(t) = \begin{cases} |t|^p/p & \text{if } |t| < T \\ T^p/p & \text{else.} \end{cases} \quad (35)$$

Computing the shrinkage rule for this mapping according to (17), we obtain

$$\nu(t) = \begin{cases} 0 & \text{if } |t| < \beta^{1/(p-2)} \\ 1 - \frac{1}{\beta}|t|^{p-2} & \text{if } \beta^{1/(p-2)} \leq |t| < T \\ 1 & \text{else,} \end{cases} \quad (36)$$

Additionally, taking $T = \infty$ we get the shrinkage rule for the unthresholded ℓ_p metric as

$$\nu(t) = \begin{cases} 0 & \text{if } |t| < \beta^{1/(p-2)} \\ 1 - \frac{1}{\beta}|t|^{p-2} & \text{else} \end{cases} \quad (37)$$




Fig. 6: Comparison of the algorithms in the presence of noise. We consider the recovery of a 256×256 original MRI brain image from its radial trajectory with 40 spokes, contaminated by Gaussian noise with standard deviation $\sigma = 18.8$. The error images are magnified by a scale of 5 fold for the best visibility. This is a challenging case due to the high undersampling factor and high measurement noise. We observe that the NLS scheme provides the best overall reconstructions.

2) *Penalty corresponding to alternating H_1 non-local scheme:* We now consider the H_1 metric, specified by

$$\varphi(t) = 1 - \exp\left(-\frac{t^2}{2\sigma^2}\right) \quad (38)$$

Computing the shrinkage rule, we obtain

$$\nu(t) = \begin{cases} 0 & \text{if } \exp\left(-\frac{t^2}{2\sigma^2}\right) > \beta\sigma^2 \\ 1 - \frac{\exp\left(-\frac{t^2}{2\sigma^2}\right)}{\beta\sigma^2} & \text{else} \end{cases} \quad (39)$$

3) *Penalty corresponding to Peyre's non-local scheme:* We now consider the penalty corresponding to Peyre's alternating scheme [8], [25]:

$$\phi(t) = 1 - \exp\left(-\frac{t}{\sigma}\right) \quad (40)$$

Computing the shrinkage rule, we obtain

$$\nu(t) = \begin{cases} 0 & \text{if } \exp\left(-\frac{t}{\sigma}\right) > \beta\sigma t \\ 1 - \frac{\exp\left(-\frac{t}{\sigma}\right)}{\beta\sigma t} & \text{else} \end{cases} \quad (41)$$

4) *Penalty corresponding to alternating non-local TV scheme:* The penalty function for the alternating non-local TV scheme is specified by [8], [26]:

$$\phi(t) = \text{erf}\left(\frac{t}{\sigma}\right) \quad (42)$$

Computing the shrinkage rule, we obtain

$$\nu(t) = \begin{cases} 0 & \text{if } \frac{2}{\sqrt{\pi}} \exp\left(-\frac{t^2}{\sigma^2}\right) > \beta\sigma t \\ 1 - \frac{2}{\sqrt{\pi}} \frac{\exp\left(-\frac{t^2}{\sigma^2}\right)}{\beta\sigma t} & \text{else} \end{cases} \quad (43)$$




Fig. 7: Comparison of different MR images using NLS algorithms in the presence of noise. We consider the recovery from a three fold undersampled radial sampling pattern, contaminated by zero mean complex Gaussian noise with standard deviation $\sigma = 10$. The top two rows show the original and reconstructed images, while the error images scale by a factor of five are shown in the bottom row. We observe that the NLS scheme preserves well the edges and the fine details at low acceleration with presence of noise.

REFERENCES

- [1] A. Buades, B. Coll, and J.M. Morel, "Denoising image sequences does not require motion estimation," in *Advanced Video and Signal Based Surveillance, 2005. AVSS 2005. IEEE Conference on*. IEEE, 2006, pp. 70–74.
- [2] S.P. Awate and R.T. Whitaker, "Unsupervised, information-theoretic, adaptive image filtering for image restoration," *IEEE Trans. Pattern Recognition*, vol. 28, pp. 364, 2006.
- [3] A. Buades, B. Coll, and J.M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling and Simulation*, vol. 4, no. 2, pp. 490–530, 2006.
- [4] L.D. Cohen, S. Bougleux, and G. Peyré, "Non-local regularization of inverse problems," in *European Conference on Computer Vision (ECCV'08)*, 2008.
- [5] G. Gilboa, J. Darbon, S. Osher, and T. Chan, "Nonlocal convex functionals for image regularization," *UCLA CAM Report*, pp. 06–57, 2006.
- [6] Y. Lou, X. Zhang, S. Osher, and A. Bertozzi, "Image recovery via nonlocal operators," *Journal of Scientific Computing*, vol. 42, no. 2, pp. 185–197, 2010.
- [7] Gabriel Peyré, Sébastien Bougleux, and Laurent Cohen, "Non-local regularization of inverse problems," in *Computer Vision—ECCV 2008*, pp. 57–68. Springer, 2008.
- [8] Zhili Yang and Mathews Jacob, "Nonlocal regularization of inverse problems: a unified variational framework," *IEEE Transactions on Image Processing*, vol. 22, no. 8, pp. 3192–3203, 2013.
- [9] Guobao Wang and Jinyi Qi, "Penalized likelihood pet image reconstruction using patch-based edge-preserving regularization," *Medical Imaging, IEEE Transactions on*, vol. 31, no. 12, pp. 2194–2204, 2012.
- [10] Z. Yang and M. Jacob, "A unified energy minimization framework for nonlocal regularization," in *IEEE ISBI*, 2011.
- [11] D. Geman and Chengda Yang, "Nonlinear image recovery with half-quadratic regularization," *IEEE Transactions on Image Processing*, vol. 4, no. 7, pp. 932–946, 1995.
- [12] Pierre Charbonnier, Laure Blanc-Féraud, Gilles Aubert, and Michel Barlaud, "Deterministic edge-preserving regularization in computed imaging," *Image Processing, IEEE Transactions on*, vol. 6, no. 2, pp. 298–311, 1997.
- [13] Alexander H. Delaney and Yoram Bresler, "Globally convergent edge-preserving regularized reconstruction: an application to limited-angle tomography," *Image Processing, IEEE Transactions on*, vol. 7, no. 2, pp. 204–221, 1998.
- [14] Mila Nikolova and Michael Ng, "Fast image reconstruction algorithms combining half-quadratic regularization and preconditioning," in *Image Processing, 2001. Proceedings. 2001 International Conference on*. IEEE, 2001, vol. 1, pp. 277–280.
- [15] R. Chartrand, "Exact reconstruction of sparse signals via nonconvex minimization," *Signal Processing Letters, IEEE*, vol. 14, no. 10, pp. 707–710, 2007.
- [16] Stanley Osher, Andrés Solé, and Luminita Vese, "Image decomposition and restoration using total variation minimization and the h¹," *Multiscale Modeling & Simulation*, vol. 1, no. 3, pp. 349–370, 2003.

- [17] Chunlin Wu, Xue-Cheng Tai, et al., "Augmented lagrangian method, dual methods, and split bregman iteration for rof, vectorial tv, and high order models.," *SIAM Journal on Imaging Sciences*, vol. 3, no. 3, pp. 300–339, 2010.
- [18] Tom Goldstein and Stanley Osher, "The split bregman method for l1-regularized problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 2, pp. 323–343, 2009.
- [19] Adrian S Lewis, "The convex analysis of unitarily invariant matrix functions," *Journal of Convex Analysis*, vol. 2, no. 1, pp. 173–183, 1995.
- [20] Joshua Trzasko and Armando Manduca, "Highly undersampled magnetic resonance image reconstruction via homotopic-minimization," *Medical imaging, IEEE Transactions on*, vol. 28, no. 1, pp. 106–121, 2009.
- [21] Yilun Wang, Junfeng Yang, Wotao Yin, and Yin Zhang, "A new alternating minimization algorithm for total variation image reconstruction," *SIAM Journal on Imaging Sciences*, vol. 1, no. 3, pp. 248–272, 2008.
- [22] <http://www.accelereyes.com>, ",".
- [23] Saiprasad Ravishankar and Yoram Bresler, "Mr image reconstruction from highly undersampled k-space data by dictionary learning," *Medical Imaging, IEEE Transactions on*, vol. 30, no. 5, pp. 1028–1041, 2011.
- [24] Michael Lustig, David Donoho, and John M Pauly, "Sparse mri: The application of compressed sensing for rapid mr imaging," *Magnetic resonance in medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [25] G. Peyré, S. Bogleux, and L.D. Cohen, "Non-local regularization of inverse problems," *Inverse Problems and Imaging*, pp. 511–530, 2011.
- [26] Yifei Lou, Xiaoqun Zhang, Stanley Osher, and Andrea Bertozzi, "Image recovery via nonlocal operators," *Journal of Scientific Computing*, vol. 42, no. 2, pp. 185–197, 2010.